

\$2.50

52

Popular Computing

July 1977 Volume 5 Number 7



How High the Precision?

A new quarterly, CRYPTOLOGIA, is published at Albion College, Albion, Michigan 49224, starting with the January 1977 issue. Individual copies are \$5; a subscription is \$16 per year.

The new journal will cover all aspects of cryptology. The initial issue includes:

"Cracking" a Random Number Generator (James Reeds)

The Biggest Bibliography (David Kahn)

Cipher Equipment (Louis Kruh)

□ □ ■ ■
□ □ ■ ■

Educational Calculator Devices (Box 974, Laguna Beach, California 92652) manufacturers a classroom device, similar to a small lecturn, that contains a Hewlett-Packard calculator, and displays in large red characters the same information as is displayed by the calculator. Thus, a lecturer can perform calculations and have all results clearly visible to a class. The price of the device depends on which HP machine is included; with the HP-25, for example, the price is \$900. (The HP-25 currently sells for \$116 alone in the Los Angeles area.)

□ □ ■ ■
□ □ ■ ■

The following cube roots:

56, 59, 60, 70, 86, 89, 91, 92, 100

(as given in issue 51, page 12) are each too high by one unit in the 18th decimal place.

(noted by H. P. Robinson)

□ □ ■ ■
□ □ ■ ■

POPULAR COMPUTING is published monthly at Box 272, Calabasas, California 91302. Subscription rate in the United States is \$20.50 per year, or \$17.50 if remittance accompanies the order. For Canada and Mexico, add \$1.50 per year to the above rates. For all other countries, add \$3.50 per year to the above rates. Back issues \$2.50 each. Copyright 1977 by POPULAR COMPUTING.

Editor: Audrey Gruenberger
Publisher: Fred Gruenberger
Associate Editors: David Babcock
Irwin Greenwald

Contributing editors: Richard Andree
William C. McGee
Thomas R. Parkin

Advertising Manager: Ken W. Sims
Art Director: John G. Scott
Business Manager: Ben Moore

Reproduction by any means is prohibited by law and is unfair to other subscribers.

How High the Precision?

Suppose you have the choice of buying one of two calculators, similar in every respect (the same function buttons, the same logic, and comparable programming capability) except for their precision level. Machine A performs all its calculations to 10-digit precision, and machine B to 12-digit precision. Machine B should cost more; perhaps 5% more. To some extent, machine B will also cost you in time; its calculations will take longer to be performed, due to the extra precision. Which machine should you buy?

The question is: what value should be placed on higher precision arithmetic? No one would seriously consider a 4-digit machine (you couldn't even balance your checkbook on it), and the 6-digit machines have largely disappeared. It seems that every customer wants, and is willing to pay for, at least 8 digits on his calculator. It also seems that, for serious work (i.e., for those who use logs and trig functions), 10 digits is the least that one will accept. Beyond that, when is even higher precision needed?

Thirty years ago, R. C. Archibald addressed that question, and answered it in effect by noting that one needs high precision for the elementary functions in order to achieve lesser precision for the higher functions.

High precision arithmetic is needed in the following situations:

- When more digits must be carried through intermediate calculations simply to get more correct digits in the result.
- When more digits must be carried through intermediate calculations in order to get a few (or any) correct digits in the result.
- When more digits must be used in order to keep a procedure from becoming unstable.
- When more digits must be carried in order to get to the correct result at all (as opposed to getting an incorrect answer to any degree of precision).
- When the problem involves numbers that are intrinsically of high precision.

When asked for problem situations in which high precision might be needed, everyone's first reaction is "Take a matrix inversion problem in which the determinant of the matrix is nearly zero." One classic case of this is Prof. Richard Andree's parallel lines problem in his book Computer Programming and Related Mathematics:

$$\left. \begin{array}{l} 3x + y = 7 \\ 2.999x + y = 1 \end{array} \right\} \quad \text{for which:} \quad \begin{array}{l} x = 6000 \\ y = -17993 \end{array}$$

and a small change in one of the coefficients:

$$\left. \begin{array}{l} 3x + y = 7 \\ 3.001x + y = 1 \end{array} \right\} \quad \text{produces:} \quad \begin{array}{l} x = -6000 \\ y = 18007 \end{array}$$

A similar example is that of Kahan's:


$$\left. \begin{array}{l} .2161x + .1441y - .1440 = 0 \\ 1.2969x + .8648y - .8642 = 0 \end{array} \right\}$$

which has the exact solution $x = 2$, $y = -2$, but the values $x = .9911$, $y = -.4870$ satisfy the equations to within 10 to the minus 8.


Both of the above examples involve two equations in two variables, so the lack of independence in the equations is fairly obvious. When the number of equations is larger than two, as in:

$$\left. \begin{array}{l} 1.23x - 3.56y + 17.85z = -99.80 \\ -6.11x + 8.40y + 12.05z = -44.98 \\ 15.91x - 27.48y + 29.44z = -209.39 \end{array} \right\}$$

it is not immediately obvious that one equation is almost a linear combination of the other two. The value of the determinant of the matrix of coefficients (.114) provides a clue, but evaluating determinants is itself a complicated calculation.



Another classic case is Wilkinson's equation (see PC23-5) in which a tiny change in one of the coefficients of a 20th degree equation changes 10 of the roots from real to imaginary.



In general, one would expect that higher precision in intermediate operations would lead to higher precision in the desired answers. But one might also expect that in some simple situations, a change in the precision level could lead to radically different results. The Error Amplification problem (PC32-1) was devised to demonstrate the latter notion. The expression:

$$2 \begin{array}{cccccccc} 3 & 5 & 7 & 9 & 11 & 13 & \dots & 99 \\ \hline 2 & 4 & 6 & 8 & 10 & 12 & \dots & 98 \end{array} \quad (EA)$$

was to be computed sequentially, alternating powers and roots, using different algorithms for the powers and roots, and it was thought that different significance levels might lead to different results. This turns out not to be so.

Dr. Mordecai Schwartz points out that the HP-67 calculator is well suited to perform research on this particular problem, since the machine can be programmed to round the number in the X (visible) register and perform subsequent calculations using the rounded number. He says:

"I constructed such a program for the EA expression, looping the y^x function 98 times (and inverting each even exponent). The output is as follows:

digital precision	calculated function value
1	1.
2	69000000.
3	30700.
4	<u>220.2</u>
5	<u>251.01</u>
6	<u>248.904</u>
7	<u>248.7874</u>
8	<u>248.81532</u>
9	<u>248.814152</u>
10	<u>248.8139420</u>
"True" value, to 10 significant digits (see PC32-1)	} ————— 248.8139802

"The underlined digits indicate accuracy to within one unit in the decimal place of the last underlined digit. The remaining digits then are 'guard digits.' Note that in 4 to 10 place precision, the 98 exponentiations (and 49 inversions) require 3 guard digits to assure accuracy to within one unit in the last desired digit. For precision to 1, 2, or 3 digits, we have insufficient guard digits and, as a result, not only does the calculated value contain no correct significant digits, but the order of magnitude is also incorrect.

"It is important to note that continued repetition of an operation does not necessarily result in a linear, a predictable, or even a continued loss of precision.

"I extended the EA expression (with the 49 even factors in the numerator and the 49 odd factors in the denominator) by programming 98 further repetitive exponentiations. The true value of the new expression is, of course, 2. The calculated value (10 digit precision) is 1.999999961. The additional 98 exponentiations (and 49 inversions) have actually yielded an increased accuracy!


"The result is accurate to the 8th digit; the remaining 2 digits are guard digits. This contrasts, as we have seen, with 7 digit accuracy and 3 guard digits for the original EA expression with half the number of exponentiations!"



Problems that intrinsically involve large numbers are relatively rare; they are frequently number theoretical problems. An example is Kraitchik's automorphic number:

$$\dots 19977392256259918212890625 = K$$

for which K^2 reproduces the same number in the low order digits (see PC19-8). Similar high-precision problems are these: the search for the largest known prime; any problem involving large factorials; research into the properties of the Fibonacci sequence; and so on. For such problems, one needs extremely high precision. The choice is not between 10 or 12 digits; one must be able to manipulate numbers with thousands of digits.



Or, consider the following puzzle, taken from Albert Beiler's Recreations in the Theory of Numbers--The Queen of Mathematics Entertains (Dover Publications T1096, 1964) problem 38 in chapter XXV:


A man had over 100 pennies, their number being a square. He divided them into 19 exactly equal square piles and had 81 pennies left over. What was the least number of pennies he had?

He could have had 11025 pennies ($= 105^2$) which, after taking away 81, would leave 19 piles of 576 each. But this is not the smallest number. The reader is encouraged to write a program to find the smallest solution, and some of the other solutions.

Usually, the computational troubles that demand higher precision are caused by the subtraction of nearly-equal numbers. For example, if a calculation includes the factor $e^x - 1$, with x very small, there can be a severe loss of precision if a direct approach is taken. If x is .00001, direct calculation (to 12 digits) gives .00001. If, however, we evaluate the series:

$$e^x = \textcircled{1} + x + \frac{x^2}{2} + \frac{x^3}{6} + \frac{x^4}{24} + \frac{x^5}{120} + \dots$$

with $x = .00001$, we obtain .000010000005.



A similar problem arises in finding the roots of the quadratic:

$$x^2 + 80x + 1 = 0.$$

The straightforward application of the quadratic formula (using 8-digit precision) gives the roots


$$x_1 = -.012502$$

$$x_2 = -79.987498$$

with a severe loss of precision for the smaller root because of the subtraction of the near-equal numbers. If the formula:

$$x = \frac{-2c}{b + \sqrt{b^2 - 4ac}}$$

is used for the smaller root (again with 8-digit precision), the result is -.01250195, which is a better approximation to the true value of -.012501953997.




Forman Acton, in Numerical Methods that Work, gives the following example. To sum the series:

$$J_0(x) = 1 - \frac{B}{(1!)^2} + \frac{B^2}{(2!)^2} - \frac{B^3}{(3!)^2} + \dots$$

with $B = (x^2/4)$. With the value of 10.1 substituted for x , "we see that the magnitudes of the terms increase up through x^{10} and then decrease steadily to insignificance. The largest absolute term is roughly 680, but it is largely offset by the negative term that follows it. We also notice that the final result, which is not large (-0.25), is reached by subtracting several large negative numbers from several large positive numbers. Since the final result has its first significant digit in the first place after the decimal point, all the figures in all the numbers to the left of that position must ultimately cancel by subtraction--and are thus irrelevant for our purposes. This means that the computed value of the largest term contains three figures that are irrelevant--three figures that occupy valuable space since they thereby displace the figures that will ultimately be needed to determine our answer. The big part of the numbers squeezes out the small, but only the small part survives in our answer. Thus, for the argument 10.1, we have barely five significant figures in our final answer if we evaluate the series term by term on a computer capable of retaining eight significant digits. Since a loss of only two significant figures is usually considered serious, we see that even for x equal to 4, the series for $J_0(x)$ is barely adequate. This behavior is typical for alternating series."

After another similar example, Acton says "The general rules are clear: Store the small quantities; compute the larger ones; never subtract nearly equal quantities."



Dr. John W. Wrench, Jr., writes "One of my favorite examples is the calculation of the periodic payment, A , required to amortize a debt of P dollars in n installments at compound interest i . The well known formula:

$$A = \frac{P \cdot i}{1 - (1+i)^{-n}}$$

reveals that we have to evaluate $(1+i)^{-n}$ to relatively high precision, and that usually requires at least eight-place logarithm tables. Here, four or five-place tables are wholly inadequate."

The following problem situation is furnished by Herman P. Robinson:

When a radio active element x disintegrates into y , and y then disintegrates, the formula is (starting with no y and $x = 1$):

$$y = \frac{\lambda_x}{\lambda_y - \lambda_x} (e^{-\lambda_x t} - e^{-\lambda_y t}),$$

where $\lambda_x = \ln 2 / (\text{half life of } x)$,


$\lambda_y = \ln 2 / (\text{half life of } y)$,

t = time from start of unit x and zero y .

If $\lambda_x = \lambda_y = \lambda$, then


$$y = \lambda t e^{-\lambda t}$$

and there is no trouble. If λ_x is almost equal to λ_y , then you are in trouble. If there are more radioactive elements in the chain, things can get worse, and you have to be very alert when calculating the yield of elements of the chain. Even present digital computers can get fouled up. Try calculating with a slide rule to get y , with $\lambda_x = 1$ and $\lambda_y = 1.01$, t various.



Consider the old puzzle problem: a chord of a circle is half a mile long and the arc it subtends is one foot longer. What is the height of the arc?

The problem is usually presented with this figure:



The figure at the right gives a more accurate picture of the height-of-the-arc problem.

The problem reduces to solving for x in:

$$\frac{\sin x}{x} = .01744 \ 66839 \ 27546 \ 49785$$

where x is half the central angle in degrees. The half angle in radians can be obtained from:

$$u = 6 \left(\frac{1}{2641} + \frac{u^2}{5!} - \frac{u^3}{7!} + \frac{u^4}{9!} - \dots \right)$$

which converges fairly rapidly; $x = \sqrt{u}$ radians. Then, with a little algebra:

$$h = \frac{2640^2 x}{5282 \left(1 + \sqrt{1 - \left(\frac{2640}{2641} x \right)^2} \right)}$$

With these results:

$$x \text{ (radians)} = .04766 \ 68105 \ 78757 \ 21801 \ 95171$$

$$x \text{ (degrees)} = 2.73110 \ 70690 \ 12333 \ 45884 \ 45373$$

$$h = 31.46605 \ 31076 \ 59484 \ 69629 \ 28678$$

Going back to Dr. Mordecai Schwartz:

"The formula $(10^6 \log 8K + 1)$ with $K = 0.1249991$ would satisfy your requirement for a relatively simple mathematical expression that yields one result if we consider K as a 4 decimal-digit number rounded to 3, and an entirely different result if we consider K as a 7 decimal-digit number rounded to 6:


$$\left[\begin{array}{l} \text{3 digit} \\ \text{precision} \end{array} \right] 10^6 \log 8(0.125) + 1 = 1.00$$

$$\left[\begin{array}{l} \text{6 digit} \\ \text{precision} \end{array} \right] 10^6 \log 8(0.124999) + 1 = -2.47437$$

"What's been done is simply to select values that yield either zero ($\log 8(0.125)$) or a number very close to zero ($\log 8(0.124999)$), and then magnify the difference by multiplication with a very large constant. Many such expressions can, of course, be constructed.

"An alternative procedure would be to construct an expression that yields either a very small positive number or a very small negative number, and again multiply by a 'magnification factor.' With such methods, we can take functions that are intrinsically continuous and create artificial 'discontinuities' based upon truncation of the operand.

"It is also possible to truncate or round arguments of discontinuous functions so that the modified arguments lie a small distance to either side of a point of discontinuity; two entirely different function values then result. Such 'magnified' or intrinsic discontinuities are not just artificial constructs. Either type could readily occur in the course of calculation if the function's behavior at critical points is overlooked."



The following problem is posed by Victor Meally of Dublin, Ireland (adapted from a problem appearing on page 111 of Mathematical Gems by Ross Honsberger).

Form the product:

$$\frac{2}{1} \cdot \frac{3}{2} \cdot \frac{5}{4} \cdot \frac{7}{6} \cdot \frac{11}{10} \cdot \frac{13}{12} \cdot \frac{17}{16} \cdot \frac{19}{18} \cdot \frac{23}{22} \dots$$

where each factor is of the form $p/(p-1)$, with the p 's successive primes.

The accompanying table shows some results. The first column is the successive primes. The second column is the product up to that prime. Column 3 shows the number of primes, K ; column 4 shows the limit just exceeded.

Mr. Meally notes that $(K-1)$, column 5, seems to form the familiar Fibonacci sequence.

The question is: is it? The product must be extended, to find when it first exceeds 10. If that occurs with the 56th prime (which is 263), then Meally's conjecture is reinforced (and we must then probe for the point at which the product first exceeds 11, and so on).

Prime	Product	Number of primes, K	Limit first exceeded	K-1
2	2.000000000	1	1	(0)
3	3.000000000	2	2	(1)
5	3.750000000	3	3	(2)
7	4.375000000	4	4	(3)
11	4.812500000	5		
13	5.213541667	6	5	(5)
17	5.539388021	7		
19	5.847131800	8		
23	6.112910518	9	6	(8)
29	6.331228751	10		
31	6.542269709	11		
37	6.723999423	12		
41	6.892099409	13		
43	7.056197013	14	7	(13)
47	7.209592601	15		
53	7.348238612	16		
59	7.474932381	17		
61	7.599514588	18		
67	7.714658748	19		
71	7.824868158	20		
73	7.933546883	21		
79	8.035259023	22	8	(21)
.....				
149	9.040126425	35	9	(34)

Still on the subject of high precision, we have the following remarks from Dr. Richard Hamming, Naval Post-graduate School, Monterey, California:

Do I Need An Accurate Computer -- Or Do I Need My Head Examined?

For example, suppose I decide to measure the size of your room by using a pair of telescopes and radar ranging to the retroreflectors on the surface of the moon. I will find that the distance I am seeking is indeed the difference between two large numbers and therefore I need an accurate computer. But just maybe I need my head examined instead!

Again, I set a pencil on its point on a slightly tilted plane, and set up the differential equations of motion with the initial conditions the angle that the pencil has at the start with two mutually perpendicular planes. I find that I need a high accuracy computer to find the subsequent motion and the final resting position. Or else I need my head examined.

I try to fit a tenth order polynomial (powers of x) to some data in a least squares sense. I find that I need a very accurate computer--or do I need my head examined? Most people who are not idiots soon realize that the powers of x are not very linearly independent, so they are not well determined by the data...but fools continue to pose idiotic problems.

Need I go on? It is true that in the field of time (or what is now the equivalent, frequency) you do have numbers known to high precision. In classic network design, the method required accurate determination of the zeros of the polynomials, even though the components of the final circuit were then described with barely three digit accuracy. Avogadro's number is given as 6.023×10^{23} but one does not need an accurate computer to use it.

[The Britannica says "The determination of this number is really an accomplishment of a high order, for it is known with more accuracy than the population of any city."]

If you owned a billion dollars, then adding or losing one cent is a change in the eleventh significant digit; who, having a billion dollars, would ever expect to know his worth to a penny? Only a fool who needed his head examined for asking the question. Could the national debt make sense to within a thousand dollars?


If you pose idiotic problems, as I have repeatedly pointed out to you, then you get idiotic requirements on the computing for those problems.

Now, regarding the selection of a hand calculator: I think the prospective purchaser had best first ask what he is going to use it for. If he will be involved with calculations leading to a determination of the speed of light (or similar problems), he needs accuracy. If his work is in general engineering, where three figures are usually rather good, all things considered, then a low accuracy (say, 7 or 8 digits) is more than he will usually need, using his hand computer. Statisticians have settled on 95% confidence limits, meaning one part in 20; even 99% confidence is just one part per 100. If the user is a number theory nut, then you know that very high precision is necessary--but he may need his head examined instead.

That mathematicians have in the past done idiotic things and given idiotic formulas is beyond argument. But that does not mean that you need high accuracy. Many of the examples you have given fail when you use your head--and if you don't, then it ought to be examined. Perhaps you should not use the series for $J_0(x)$ for large values of x , but rather find the series for $1/J_0(x)$, noting that the latter has all positive terms and thus avoids the intrinsic problems of alternating series.

I do not for the moment question the ability of anyone to find problems that seem to require high precision, but I do wonder about whether the problem is worth solving in the form given. Idiotic problems are easy to pose, and when hard to compute are probably wrongly posed.

My experience in 30 years of computing is that, usually, an apparent need for high precision means faulty formulation of the problem, and that reformulation will usually yield more dividends with less computing. These are both gains, but the former is greater.



Comments on Dr. Hamming's essay by FJG:


It is difficult to quarrel with the truisms that Hamming expresses. Many of them are like the directions in the manual for my car: "Do not bump your head on the raised trunk lid."

The trouble is, the pathological cases in numerical work do not come to us bearing labels that say "Look out--trouble ahead." For most of us, lacking 30 years of experience in numerical work, the troubles are booby traps, hidden among quite normal calculations. What is worse, a normal situation may become pathological in the middle of a long calculation, when a variable goes critical.

The quadratic equation example is a fine illustration. The old faithful quadratic formula works just fine most of the time. If I can even remember that it may on occasion trip me up, should I insert into my program a test for the pathological case, together with a re-calculation using the alternate formula? I may be in even worse trouble then, since the million normal cases involved in my problem now chew up excessive computer time.

It would seem that if the cost is not too great (e.g., 5% more for the hardware), that having extra guard digits in my calculations will give me some protection from things that can go wrong that I might not even be aware of.

We are accustomed to being protected from bad things that can occur in long calculations. We buffer our divisions with a test that says, effectively, "Am I about to attempt a division by zero?" We expect square root and logarithm subroutines to furnish courteous error indications when we offer them negative arguments. We would appreciate having our floating point subroutines tell us when we have exceeded the range of the system (we don't always get this goodie). And so it goes: I would personally put a high premium on high precision, quite apart from my proclivity for number theory problems. I want all the protection I can get.



Further comments by Associate Editor David Babcock:

I find the viewpoints of both Hamming and Gruenberger partly correct and partly incorrect.

Hamming has made the most important point; namely, that the chief criterion for selecting one machine over another is its intended use. The cost difference between 8-digit, 10-digit, and 12-digit machines is relatively low, both in time and money. The real cost difference is in the functions: how many there are, and how they are engineered for ease of use. I would like to see the JOSS (and APL) philosophy followed by the calculator manufacturers; namely, that every function should behave as you would expect it to. This doesn't necessarily mean carrying lots of extra digits, but rather more care in the algorithms (e.g., testing for magic values).

I don't go along with the idea of giving the user more digits to "buffer" him from problem areas. Should we put 5-foot bumpers on cars in case the driver is careless? This sounds to me like operating on the symptoms, rather than the disease. Granted, we can't expect new owners of calculators to be whizzes at numerical methods, but that's all the more reason to design the calculator carefully to produce correct results (or else to have it give an indication that the function has been used incorrectly).

On the other hand, I feel that there are places where high precision is necessary and that at least the tools (that is, many digits) should be available to the user. If you were the Bank of America balancing its books you can bet that they would worry about the pennies.

Hamming advocates redefining the problem so as to obviate the need for high precision. One could redefine an equation that uses logarithms (e.g., write out the series) but would Hamming advocate removing the LOG button from his calculator?--I think not.

I agree that many of the problems that have appeared in POPULAR COMPUTING, which require high precision, are artificial, but I think that Hamming misses the point of those problems. We need a large stock of simple problems (readily grasped by beginners) in order to learn how to use a new tool. The beginner can then go on to apply that new tool to his real problems. Hamming must see some need (and educational value) to high precision arithmetic, or why else did he spend his time learning it and learning to apply it?

■ ■
■ ■

All the above discussion on the need for or utility of extra precision leaves the question still open. Other examples and opinions are solicited, to be presented in subsequent issues.

More on Albert Beiler's problem (page 7)

of the 19 piles of pennies:

Beiler's problem requires solutions in integers of the equation:

$$x^2 - 19y^2 = 81.$$

The formulas on the right give an infinite number of solutions, but not all the solutions or even the smallest ones.

This is one of those situations in which the numbers involved intrinsically demand high precision. One of the factors in the formulas on the right is the number 340 raised to high powers. Even for the 4th power, the capacity of a 10-digit calculator is exceeded.

For $n = 1$, the two solutions given by the equations are:

$$959^2 - 19 \cdot 220^2 = 81$$

$$2441^2 - 19 \cdot 560^2 = 81.$$

■ ■
■ ■


$$x = 10 \left[\frac{(170 + 39\sqrt{19})^n}{2} + \frac{(170 - 39\sqrt{19})^n}{2} \right] + 19 \left[\frac{(170 + 39\sqrt{19})^n}{2\sqrt{19}} - \frac{(170 - 39\sqrt{19})^n}{2\sqrt{19}} \right]$$

$$y = \left[\frac{(170 + 39\sqrt{19})^n}{2} + \frac{(170 - 39\sqrt{19})^n}{2} \right] + 10 \left[\frac{(170 + 39\sqrt{19})^n}{2\sqrt{19}} - \frac{(170 - 39\sqrt{19})^n}{2\sqrt{19}} \right]$$

Mr. Anderson's Problem

Take a 4-digit number. Duplicate the high order digit on the right end of the number. Form the absolute differences of the adjacent digits; this produces a new 4-digit number. Repeat the process, and count the number of such steps it takes to get to 0000.

For example:



So for the number 1480, the result is 6 steps.

Of the 10,000 4-digit numbers, one (0000) obviously goes out immediately; that is, in zero steps. Nine numbers (1111, 2222, ..., 9999) go out in one step.

That leaves 9990 numbers to be accounted for. A computer solution to this problem is an instructive exercise in coding.




Take a decimal fraction and convert it to its binary version. Add these two as decimal numbers, producing a new decimal fraction. Convert that fraction to binary and subtract. Repeat the process, alternately adding and subtracting. For example, starting with $\frac{4}{7}$:

decimal:	.5714285714285
binary:	.1001001001001
the sum:	.6715286715286
binary:	.1010101111100
the difference:	.5705185604185
binary:	.1001001000001
the sum:	.6706186604186

...and so on; the above is three iterations.

For numbers between zero and about .2, this process converges back to the original number. For starting fractions greater than .2, the process oscillates, resulting in two fractions (that is, the result taken after an addition or after a subtraction).



So far, so good. It's a weird procedure, but it makes an excellent coding exercise for beginners.

If the procedure is carried out for each fraction to a uniform depth (say, 40 iterations), the plot of the results, shown here, is surprising. The changes at .5 and .75 are discontinuities, with the lines of discontinuity at .501540154 and .75155401.

This process needs further investigation.

What's Going On Here?...The Binary/Decimal Game